Refine Search

Search Results -

Terms	Documents
6014670.pn.	2

Database:
US Pre-Grant Publication Full-Text Database
US Patents Full-Text Database
US OCR Full-Text Database
EPO Abstracts Database
JPO Abstracts Database
Derwent World Patents Index
IBM Technical Disclosure Bulletins

Search:

Refine Search

Recall Text | Clear | Interrupt

Search History

DATE: Saturday, August 30, 2003 Printable Copy Create Case

Set Name side by side	Query	Hit Count	Set Name result set
DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR			
L6	6014670.pn.	2	L6
L5	5978788.pn.	2	L5
L4	"star schema" and (datawarehouse or data with warehouse or datamart or data with mart)	90	L4
L3	"reverse star schema"	7	L3
L2	"reverse star schema".ti.ab.	16	L2
L1	"reverse stare schema".ti.ab.	13	L1

END OF SEARCH HISTORY

Hit List

Clear | Generate Collection | Print | Fwd Refs | Bkwd Refs | Generate OACS

Search Results - Record(s) 1 through 7 of 7 returned.

☐ 1. Document ID: US 20030115211 A1

L3: Entry 1 of 7 File: PGPB Jun 19, 2003

PGPUB-DOCUMENT-NUMBER: 20030115211
PGPUB-FILING-TYPE: new
DOCUMENT-IDENTIFIER: US 20030115211 A1

TITLE: Spatial intelligence system and method

PUBLICATION-DATE: June 19, 2003

INVENTOR-INFORMATION:

NAME	CITY	STATE	COUNTRY	RULE-47
Chen, Li-Wen	Cupertino	CA	US	
Luu, Victor	Morgan Hill	CA	US	

US-CL-CURRENT: 707/102; 707/104.1

Full | Title | Citation | Front | Review | Classification | Date | Reference | Sequences | Attachments | Claims | KWIC | Draw De

☐ 2. Document ID: US 20020116389 A1

L3: Entry 2 of 7 File: PGPB Aug 22, 2002

PGPUB-DOCUMENT-NUMBER: 20020116389
PGPUB-FILING-TYPE: new
DOCUMENT-IDENTIFIER: US 20020116389 A1

TITLE: System and method for providing a data warehouse in accordance with a virtual schema

PUBLICATION-DATE: August 22, 2002

INVENTOR-INFORMATION:

NAME	CITY	STATE	COUNTRY	RULE-47
Chen, Li-Wen	Cupertino	CA	US	
Ortiz, Juan J.	Hayward	CA	US	

US-CL-CURRENT: 707/103R

☐ 3. Document ID: US 20020059183 A1

L3: Entry 3 of 7 File: PGPB May 16, 2002

PGPUB-DOCUMENT-NUMBER: 20020059183
PGPUB-FILING-TYPE: new
DOCUMENT-IDENTIFIER: US 20020059183 A1

TITLE: Method and system for managing event attributes

PUBLICATION-DATE: May 16, 2002

INVENTOR-INFORMATION:

NAME	CITY	STATE	COUNTRY	RULE-47
Chen, Li-Wen	Cupertino	CA	US	

US-CL-CURRENT: 707/1; 705/8, 705/9, 707/104.1

☐ 4. Document ID: US 20020038230 A1

L3: Entry 4 of 7 File: PGPB Mar 28, 2002

PGPUB-DOCUMENT-NUMBER: 20020038230
PGPUB-FILING-TYPE: new
DOCUMENT-IDENTIFIER: US 20020038230 A1

TITLE: User interface and method for analyzing customer behavior based upon event attributes

PUBLICATION-DATE: March 28, 2002

INVENTOR-INFORMATION:

NAME	CITY	STATE	COUNTRY	RULE-47
Chen, Li-Wen	Cupertino	CA	US	

US-CL-CURRENT: 705/7

☐ 5. Document ID: US 6411961 B1

L3: Entry 5 of 7 File: USPT Jun 25, 2002

US-PAT-NO: 6411961
DOCUMENT-IDENTIFIER: US 6411961 B1

TITLE: Apparatus for providing a reverse star schema data model

DATE-ISSUED: June 25, 2002

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Chen; Li-Wen	Cupertino	CA		

US-CL-CURRENT: 707/102; 705/10, 707/104.1

Full | Title | Citation | Front | Review | Classification | Date | Reference | Sequences | Attachments | Claims | KWIC | Draw De

---

☐ 6. Document ID: US 6377934 B1

L3: Entry 6 of 7 File: USPT Apr 23, 2002

US-PAT-NO: 6377934
DOCUMENT-IDENTIFIER: US 6377934 B1

TITLE: Method for providing a reverse star schema data model

DATE-ISSUED: April 23, 2002

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Chen; Li-Wen	Cupertino	CA		
Ortiz; Juan J.	Hayward	CA		

US-CL-CURRENT: 705/10; 707/102, 707/104.1

Full | Title | Citation | Front | Review | Classification | Date | Reference | Sequences | Attachments | Claims | KWIC | Draw De

---

☐ 7. Document ID: US 6363353 B1

L3: Entry 7 of 7 File: USPT Mar 26, 2002

US-PAT-NO: 6363353
DOCUMENT-IDENTIFIER: US 6363353 B1

TITLE: System for providing a reverse star schema data model

DATE-ISSUED: March 26, 2002

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Chen; Li-Wen	Cupertino	CA		

US-CL-CURRENT: 705/10; 707/102, 707/104.1

Full | Title | Citation | Front | Review | Classification | Date | Reference | Sequences | Attachments | Claims | KWIC | Draw De

Clear | Generate Collection | Print | Fwd Refs | Bkwd Refs | Generate OACS

Terms	Documents
"reverse star schema"	7

Display Format: - Change Format

Previous Page Next Page Go to Doc#

Refine Search

Search Results -

Terms	Documents
L7 and (rolap or metacube same data near warehouse)	55

Database:
US Pre-Grant Publication Full-Text Database
US Patents Full-Text Database
US OCR Full-Text Database
EPO Abstracts Database
JPO Abstracts Database
Derwent World Patents Index
IBM Technical Disclosure Bulletins

Search:

Refine Search

Recall Text

Clear

Interrupt

Search History

DATE: Saturday, August 30, 2003 Printable Copy Create Case

Set Name side by side	Query	Hit Count	Set Name result set
DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR			
L18	L7 and (rolap or metacube same data near warehouse)	55	L18
L17	L16 and customer near profile	3	L17
L16	L15 and fact near table	39	L16
L15	l7 and dimension$ near (database or data with base)	93	L15
L14	L13 and fact near table	12	L14
L13	L12 and hierarc$	73	L13
L12	l7 and customer near profile	129	L12
DB=USPT; PLUR=YES; OP=OR			
L11	6004276.pn.	1	L11
L10	6325540.pn.	1	L10
DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR			
L9	L8 and customer with profile	11	L9
L8	L7 and dimension near table	118	L8

L7	(data with warehouse or datawarehouse or data with mart or datamart)	3309	L7
L6	L5 and (data with warehouse or data with mart)	0	L6
L5	5701400.pn.	2	L5
L4	L1 and (data with warehouse or data with mart)	0	L4
L3	L1 and data with warehouse or data with mart	272	L3
L2	L1 and (datawarehouse or data with warehouse or datamart or data with mart)	0	L2
L1	5822744.pn.	2	L1

END OF SEARCH HISTORY

Search Results

Search Results for: **[rolap <and> aggregate <and> views]**
Found **47** of **120,890 searched.**

Search within Results

GO > Advanced Search

> Search Help/Tips

Sort by: Title Publication Publication Date Score Binder

Results 41 - 47 of 47 short listing

Prev Page 1 2 3 Next Page

41 Requirement-based data cube schema design 77%
David W. Cheung , Bo Zhou , Ben Kao , Hongjun Lu , Tak Wah Lam , Hing Fung Ting
Proceedings of the eighth international conference on Information and knowledge management November 1999
On-line analytical processing (OLAP) requires efficient processing of complex decision support queries over very large databases. It is well accepted that pre-computed data cubes can help reduce the response time of such queries dramatically. A very important design issue of an efficient OLAP system is therefore the choice of the right data cubes to materialize. We call this problem the data cube schema design problem. In this paper we show that the problem of finding an op ...

42 Compressed data cubes for OLAP aggregate query approximation on continuous dimensions 77%
Jayavel Shanmugasundaram , Usama Fayyad , P. S. Bradley
Proceedings of the fifth ACM SIGKDD international conference on Knowledge discovery and data mining August 1999

43 Simultaneous optimization and evaluation of multiple dimensional queries 77%
Yihong Zhao , Prasad M. Deshpande , Jeffrey F. Naughton , Amit Shukla
ACM SIGMOD Record , Proceedings of the 1998 ACM SIGMOD international conference on Management of data June 1998
Volume 27 Issue 2
Database researchers have made significant progress on several research issues related to multidimensional data analysis, including the development of fast cubing algorithms, efficient schemes for creating and maintaining precomputed group-bys, and the design of efficient storage structures for multidimensional data. However, to date there has been little or no work on multidimensional query optimization. Recently, Microsoft has proposed “OLE DB

for OLAP” as a standard multidime ...

44 Dynamic assembly of views in data cubes 77%
John R. Smith , Vittorio Castelli , Anant Jhingran , Chung-Sheng Li
Proceedings of the seventeenth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems May 1998

45 The IBM data warehouse architecture 77%
Charles Bontempo , George Zagelow
Communications of the ACM September 1998
Volume 41 Issue 9

46 Quasi-cubes: exploiting approximations in multidimensional databases 77%
Daniel Barbará , Mark Sullivan
ACM SIGMOD Record September 1997
Volume 26 Issue 3
A data cube is a popular organization for summary data. A cube is simply a multidimensional structure that contains at each point an aggregate value, i.e., the result of applying an aggregate function to an underlying relation. In practical situations, cubes can require a large amount of storage. The typical approach to reducing storage cost is to materialize parts of the cube on demand. Unfortunately, this lazy evaluation can be a time-consuming operation. In this paper, we desc ...

47 Data warehousing and OLAP for decision support 77%
Surajit Chaudhuri , Umeshwar Dayal
ACM SIGMOD Record , Proceedings of the 1997 ACM SIGMOD international conference on Management of data June 1997
Volume 26 Issue 2
On-Line Analytical Processing (OLAP) and Data Warehousing are decision support technologies. Their goal is to enable enterprises to gain competitive advantage by exploiting the ever-growing amount of data that is collected and stored in corporate databases and files for better and faster decision making. Over the past few years, these technologies have experienced explosive growth, both in the number of products and services offered, and in the extent of coverage in the tra ...

Results 41 - 47 of 47 short listing

Prev Page 1 2 3 Next Page

Nothing Found

Your search for the *Phrase* **maintainance <and> data <and> cubes <and> summary <and> tables<AND>((rolap <and> aggregate <and> views))** did not return any results.
To search for *terms* separate them with **AND** or **OR**.
Click on the suggested options:

maintainance AND <and> AND data AND <and> AND cubes AND <and> AND summary AND <and> AND tables<AND>((rolap AND <and> AND aggregate AND <and> AND views) AND)

maintainance OR OR data OR OR cubes OR OR summary OR OR tables((rolap OR OR aggregate OR OR views) OR)

To search for names try using only the last or first name.

You may revise it and try your search again below or click advanced search for more options.

```
maintainance <and> data <and> cubes
<and> summary <and>
tables<AND>((rolap <and> aggregate
<and> views) )
```

SEARCH [Advanced Search] [Search Help/Tips]

Complete Search Help and Tips

The following characters have specialized meaning:

Special Characters	Description
, ( ) [	These characters end a text token.
= > < !	These characters end a text token because they signify the start of a field operator. (! is special: != ends a token.)
` @ \Q < { [!	These characters signify the start of a delimited token. These are terminated by the end character associated with the start character.

L18: Entry 32 of 55 File: USPT Apr 15, 2003

DOCUMENT-IDENTIFIER: US 6549907 B1
TITLE: Multi-dimensional database and data cube compression for aggregate query support on numeric dimensions

Brief Summary Text (9):
As data warehousing becomes more popular, OLAP is gaining in importance as a primary interface to evaluating data contained in the data warehouse. Most successful data mining applications include reporting systems having fast query response mechanisms. Most corporations require decision support and would benefit from improved technology to help in making decisions based upon rapidly gathered and organized data.

Other Reference Publication (33):
Y. Kotidis, N. Rousopoulos, "An Alternative Storage Organization for ROLAP Aggregate Views Based On Cubetrees", Proc., ACM SIGMOD Conf., Seattle, 1998.

Other Reference Publication (34):
I. Mumick et al., "Maintainance of Data Cubes and Summary Tables in a Warehouse", Proc. ACM SIGMOD Conf., Tuscon, 1997.

Data Warehousing and OLAP for Decision Support

Surajit Chaudhuri
Microsoft Research, Redmond
surajitc@microsoft.com

Umeshwar Dayal
Hewlett-Packard Laboratories, Palo Alto
dayal@hpl.hp.com

Description

On-Line Analytical Processing (OLAP) and Data Warehousing are *decision support* technologies. Their goal is to enable enterprises to gain competitive advantage by exploiting the ever-growing amount of data that is collected and stored in corporate databases and files for better and faster decision making. Over the past few years, these technologies have experienced explosive growth, both in the number of products and services offered, and in the extent of coverage in the trade press. Vendors, including all database companies, are paying increasing attention to all aspects of decision support.

Decision support places some rather different requirements on database technology as compared to traditional on-line transaction processing (OLTP) applications. OLTP applications typically automate clerical data processing tasks such as order entry and banking transactions that are the bread-and-butter, day-to-day operations of an organization. These tasks are structured and repetitive, and consist of short, atomic, isolated transactions, which require detailed, up-to-date data, and read or update a few (tens of) records. Consistency and recoverability of the database are critical, and maximising transaction throughput is the key performance metric.

Decision support, in contrast, requires historical, summarised and consolidated data from many sources scattered through the enterprise. Data is extracted from these sources and loaded into a data warehouse, a large, integrated, relatively static, database that is often maintained separately from the organization's operational databases. To facilitate complex analyses and visualisation, the data warehouse typically supports a *multidimensional* model of data. Since data warehouses contain consolidated data, perhaps from several operational databases, over potentially long periods of time, they tend to be orders of magnitude larger than operational databases; enterprise data warehouses are projected to be hundreds of gigabytes to terabytes in size. The workloads are query intensive with mostly *ad hoc*, complex queries that can access millions of records. Query throughput and response times are more important than transaction throughput.

SIGMOD '97 AZ,USA

Data warehouses might be implemented on standard or extended relational database management systems, called *Relational OLAP (ROLAP)* servers. These servers assume that data is stored in relational databases, using special database designs (star and snowflake schemas) to represent the multidimensional data model; special access methods and query processing techniques to efficiently map OLAP operations on the underlying relational database. Alternatively, *multidimensional OLAP (MOLAP)* servers may be used. These are specialized servers that directly store multidimensional data in special data structures (e.g., arrays) and implement the OLAP operations over these special data structures.

This tutorial provides a roadmap of data warehousing and OLAP technologies, with an emphasis on their new requirements. We describe back end tools for extracting, cleaning and loading data into a data warehouse; multidimensional data models and OLAP operations; front end client tools for querying and data analysis; server extensions for efficient query processing; and tools for metadata management and for managing the warehouse. We survey the state of the art and mention representative products. In a recent overview paper, we have summarised the issues that are discussed in this tutorial [1].

The area opens up interesting research directions, with ties to past work in database systems, but with different assumptions and requirements. Only very recently, however, has the database research community started to address some of these issues. Research in data warehousing so far has focused primarily on query processing and view maintenance issues. There still are many open research problems. We describe some of these briefly.

Outline

1. Introduction
   - definitions, evolution, differences from OLTP, architectures
2. Models and Tools
   - conceptual model for OLAP
   - front-end tools (e.g., multidimensional spreadsheets)
   - database design (e.g., star and snowflake schema)
3. Database Server technologies for Decision Support Queries
   - specialized indexing and query processing techniques

- intelligent processing of aggregates
- complex query processing
- extensions to SQL
- ROLAP vs MOLAP

4. Other Services for OLAP/Data warehousing
   - data cleaning, loading and refresh
   - tools for warehouse, system and process management
   - metadata management and the role of repository
5. State of Commercial Practice
6. Research Issues

References

[1] S. Chaudhuri, and U. Dayal, "An Overview of Data Warehousing and OLAP Technology", SIGMOD Record, March 1997.

Citation

International Conference on Management of Data and Symposium on Principles of Database Systems >archive
Proceedings of the 1997 ACM SIGMOD international conference on Management of data >toc
1997 , Tucson, Arizona, United States

Data warehousing and OLAP for decision support

> Also published in ...

Authors
Surajit Chaudhuri
Umeshwar Dayal

Sponsor
SIGMOD : ACM Special Interest Group on Management of Data

Publisher
ACM Press New York, NY, USA

Pages: 507 - 508 Series-Proceeding-Article
Year of Publication: 1997
ISBN:0-89791-911-4

doi> http://doi.acm.org/10.1145/253260.253373 (Use this link to Bookmark this page)

> full text > abstract > references > citings > index terms > peer to peer

> Discuss > Similar > Review this Article Save to Binder

> BibTex Format

↑ FULL TEXT: Access Rules

pdf 223 KB

↑ ABSTRACT

On-Line Analytical Processing (OLAP) and Data Warehousing are decision support technologies. Their goal is to enable enterprises to gain competitive advantage by exploiting the ever-growing amount of data that is collected and stored in corporate databases and files for better and faster decision making. Over the past few years, these technologies have experienced explosive growth, both in the number of products and services offered, and in the extent of coverage in the trade press. Vendors, including all

database companies, are paying increasing attention to all aspects of decision support.

REFERENCES

Note: OCR errors may be found in this Reference List extracted from the full text article. ACM has opted to expose the complete List rather than only correct and linked references.

1 Surajit Chaudhuri , Umeshwar Dayal, An overview of data warehousing and OLAP technology, ACM SIGMOD Record, v.26 n.1, p.65-74, March 1997

CITINGS *2*

Lixin Fu , Joachim Hammer, CubiST: a new algorithm for improving the performance of ad-hoc OLAP queries, Proceedings of the third ACM international workshop on Data warehousing and OLAP, p.72-79, November 06-11, 2000, McLean, Virginia, United States

Stefan Berchtold , Christian Böhm , Hans-Peter Kriegal, The pyramid-technique: towards breaking the curse of dimensionality, ACM SIGMOD Record, v.27 n.2, p.142-153, June 1998

INDEX TERMS

Primary Classification:
H. Information Systems
↳ **H.4** INFORMATION SYSTEMS APPLICATIONS
↳ **H.4.2** Types of Systems
↳ **Subjects:** Decision support (e.g., MIS)

Additional Classification:
H. Information Systems
↳ **H.2** DATABASE MANAGEMENT
↳ **H.3** INFORMATION STORAGE AND RETRIEVAL

General Terms:
Design, Management, Performance, Theory

Peer to Peer - Readers of this Article have also read:

- Data structures for quadtree approximation and compression
  Communications of the ACM 28, 9
  Hanan Samet

- The state of the art in automating usability evaluation of user interfaces
  ACM Computing Surveys (CSUR) 33, 4

- A lifecycle process for the effective reuse of commercial off-the-shelf (COTS) software
  Proceedings of the 1999 symposium on Software reusability
  Christine L. Braun

- Using adapters to reduce interaction complexity in reusable component-based software development
  Proceedings of the 1999 symposium on Software reusability
  David Rine , Nader Nada , Khaled Jaber

- A catalog of techniques for resolving packaging mismatch
  Proceedings of the 1999 symposium on Software reusability
  Robert DeLine

↑ This Article has also been published in:

- ACM SIGMOD Record
  Volume 26 , Issue 2 (June 1997)
